Interference Searched

EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	1	(cache with (RAID near5 configuration)).clm.	US-PGPUB	OR	OFF	2006/06/20 12:49
L2	1	(cache with (RAID near5 (reconfigurat$3 or configurat$3))).clm.	US-PGPUB	OR	OFF	2006/06/20 12:55
L3	3	(cache same (RAID near5 (reconfigurat$3 or configurat$3))).clm.	US-PGPUB	OR	OFF	2006/06/20 12:56
L4	3	(cache same (RAID near5 (reconfigurat$3 or configurat$3 or (chang$3 near2 level)))).clm.	US-PGPUB	OR	OFF	2006/06/20 12:56
L5	3	(cach$3 same (RAID near5 (reconfigurat$3 or configurat$3 or (chang$3 near2 level)))).clm.	US-PGPUB	OR	OFF	2006/06/20 12:59
L6	0	(cach$3 same ((chang$3 or alter$3 or updat$3 or modify$3) near5 (RAID near2 level))).clm.	US-PGPUB	OR	OFF	2006/06/20 13:00
L7	1	(cach$3 and ((chang$3 or alter$3 or updat$3 or modify$3) near5 (RAID near2 level))).clm.	US-PGPUB	OR	OFF	2006/06/20 13:00

EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L2	1706	(RAID or (redundant near2 array near2 inexpensive near2 disk$1)) same cache	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 10:40
L3	358	(read$3 or load$3) with (RAID or (redundant near2 array near2 inexpensive near2 disk$1)) same cache	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 11:11
L4	238	(read$3 or load$3) with (RAID or (redundant near2 array near2 inexpensive near2 disk$1)) with cache	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 10:01
L5	11	(read$3 or load$3) with cache with ((RAID or (redundant near2 array near2 inexpensive near2 disk$1)) near5 configuration)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 10:07
L6	47	cache with ((RAID or (redundant near2 array near2 inexpensive near2 disk$1)) near5 configuration)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 10:09
L7	70	cache with (((RAID or (redundant near2 array near2 inexpensive near2 disk$1)) near5 (information or configuration)))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 10:11
L8	13	cache with (RAID or (redundant near2 array near2 inexpensive near2 disk$1)) with (information or configuration) with (RAId adj level)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 10:13
L9	13	cache with (RAID or (redundant near2 array near2 inexpensive near2 disk$1)) with (information or configuration) with (RAId near2 level)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 10:14

EAST Search History

L10	13	cache with (RAID or (redundant near2 array near2 inexpensive near2 disk$1)) with (information or configuration) with (RAID near2 level)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 10:31
L11	4	((read$3 or load$3) near5 cache) with (RAID or (redundant near2 array near2 inexpensive near2 disk$1)) with (information or configuration) with (RAID near2 level)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 10:39
L12	1278	(711/114).CCLS.	USPAT; USOCR	OR	OFF	2006/06/20 10:39
L13	1604	(RAID or (redundant near2 array near2 inexpensive near2 disk$1)) near2 level	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 10:41
L14	128	((RAID or (redundant near2 array near2 inexpensive near2 disk$1)) near2 level) near10 (chang$3 or modify$3 or alter$3)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 11:08
L15	49	((RAID or (redundant near2 array near2 inexpensive near2 disk$1)) near2 level) near10 (chang$3 or modify$3 or alter$3) near10 (configuration or information)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 10:43
L16	2	(((RAID or (redundant near2 array near2 inexpensive near2 disk$1)) near2 level) near10 (chang$3 or modify$3 or alter$3) near10 (configuration or information)) same cache	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 10:47
L17	14	(((RAID or (redundant near2 array near2 inexpensive near2 disk$1)) near2 level) near10 (chang$3 or modify$3 or alter$3 or updat$3)) same cache	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 10:51
L18	27	(RAID or (redundant near2 array near2 inexpensive near2 disk$1)) near5 (reconfiguration or re-configuration)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 11:06

EAST Search History

L19	1	((RAID or (redundant near2 array near2 inexpensive near2 disk$1)) near5 (reconfiguration or re-configuration)) same cache	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 10:53
L20	1379	(711/118).CCLS.	USPAT; USOCR	OR	OFF	2006/06/20 10:53
L21	1487	(711/154).CCLS.	USPAT; USOCR	OR	OFF	2006/06/20 10:53
L22	809	(711/113).CCLS.	USPAT; USOCR	OR	OFF	2006/06/20 11:06
L23	1436	(714/6).CCLS.	USPAT; USOCR	OR	OFF	2006/06/20 11:07
L24	531	(714/7).CCLS.	USPAT; USOCR	OR	OFF	2006/06/20 11:07
L25	1336	(711/170).CCLS.	USPAT; USOCR	OR	OFF	2006/06/20 11:07
L26	0	14 and 20	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 11:08
L27	2	14 and 21	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 11:08
L28	5	14 and 22	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 11:08
L29	25	14 and 23	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 11:08
L30	10	14 and 24	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 11:08

EAST Search History

L31	9	14 and 25	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 11:08
L32	8	3 and 21	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 11:09
L33	34	3 and 22	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 11:09
L34	40	3 and 23	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 11:09
L35	10	3 and 24	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 11:09
L36	12	3 and 25	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 11:09
L37	6	(read$3 or load$3) with (RAID or (redundant near2 array near2 inexpensive near2 disk$1)) same cache same (active or on)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 11:12
L38	6	(read$3 or load$3) with (RAID or (redundant near2 array near2 inexpensive near2 disk$1)) same cache same (active)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/06/20 11:12

Results for **"(raid reconfiguration<and>cache)"**
Your search matched **3** of **1360403** documents.
A maximum of **100** results are displayed, **25** to a page, sorted by **Relevance** in **Descending** order.

e-mail

» Search Options

View Session History

New Search

Modify Search

(raid reconfiguration<and>cache) Search

Check to search only within this results set

Display Format: Citation Citation & Abstract

» Key

IEEE JNL	IEEE Journal or Magazine
IEE JNL	IEE Journal or Magazine
IEEE CNF	IEEE Conference Proceeding
IEE CNF	IEE Conference Proceeding
IEEE STD	IEEE Standard

view selected items **Select All** **Deselect All**

1. **Reliability and performance of RAIDs**
Schwarz, T.J.E.; Burkhard, W.A.;
Magnetics, IEEE Transactions on
Volume 31, Issue 2, Mar 1995 Page(s):1161 - 1166
Digital Object Identifier 10.1109/20.364801
AbstractPlus | Full Text: PDF(556 KB) **IEEE JNL**
Rights and Permissions

2. **5th Annual Magnetic Recording Conference (TMRC) on Signal Processin**
Magnetics, IEEE Transactions on
Volume 31, Issue 2, Mar 1995
Digital Object Identifier 10.1109/20.364762
AbstractPlus | Full Text: PDF(14616 KB) **IEEE JNL**
Rights and Permissions

3. **Proceedings of the Seventh Euromicro Workshop on Parallel and Distrib PDP'99**
Parallel and Distributed Processing, 1999. PDP '99. Proceedings of the Seven Workshop on
3-5 Feb. 1999
Digital Object Identifier 10.1109/EMPDP.1999.746637
AbstractPlus | Full Text: PDF(208 KB) **IEEE CNF**
Rights and Permissions

Results for **"(change raid level<and>cache)"**
Your search matched **0** documents.
A maximum of **100** results are displayed, **25** to a page, sorted by **Relevance** in **Descending** order.

e-mail

» Search Options

View Session History

New Search

Modify Search

(change raid level<and>cache) Search

Check to search only within this results set

Display Format: Citation Citation & Abstract

» Key

IEEE JNL	IEEE Journal or Magazine
IEE JNL	IEE Journal or Magazine
IEEE CNF	IEEE Conference Proceeding
IEE CNF	IEE Conference Proceeding
IEEE STD	IEEE Standard

No results were found.

Please edit your search criteria and try again. Refer to the Help pages if you need assistan search.

Results for **"(raid configuration<and>cache)"** e-mail
Your search matched **13** of **1360403** documents.
A maximum of **100** results are displayed, **25** to a page, sorted by **Relevance** in **Descending** order.

» Search Options

View Session History

New Search

» Key

IEEE JNL	IEEE Journal or Magazine
IEE JNL	IEE Journal or Magazine
IEEE CNF	IEEE Conference Proceeding
IEE CNF	IEE Conference Proceeding
IEEE STD	IEEE Standard

Modify Search

(raid configuration<and>cache) Search

Check to search only within this results set

Display Format: Citation Citation & Abstract

view selected items **Select All** **Deselect All**

1. **Interplay of energy and performance for disk arrays running transaction | workloads**
Gurumurthi, S.; Jianyong Zhang; Sivasubramaniam, A.; Kandemir, M.; Franke,
Vijaykrishnan, N.; Irwin, M.J.;
Performance Analysis of Systems and Software, 2003. ISPASS. 2003 IEEE Int Symposium on
6-8 March 2003 Page(s):123 - 132
AbstractPlus | Full Text: PDF(922 KB) **IEEE CNF**
Rights and Permissions

2. **The RAID configuration tool**
Zabback, P.; Menon, J.; Riegel, J.;
High Performance Computing, 1996. Proceedings. 3rd International Conferenc
19-22 Dec. 1996 Page(s):55 - 61
Digital Object Identifier 10.1109/HIPC.1996.565797
AbstractPlus | Full Text: PDF(536 KB) **IEEE CNF**
Rights and Permissions

3. **Toward the age of smarter storage**
Robinson, G.S.;
Computer
Volume 35, Issue 12, Dec. 2002 Page(s):35 - 41
Digital Object Identifier 10.1109/MC.2002.1106177
AbstractPlus | Full Text: PDF(359 KB) **IEEE JNL**
Rights and Permissions

4. **Channel system requirements and costs**
Goldsmith, B.J.; Adams, R.C.N.;
Broadcasting Convention, 1997. International
12-16 Sept. 1997 Page(s):246 - 251
AbstractPlus | Full Text: PDF(508 KB) **IEE CNF**

5. **RAID keeps going and going and... [magnetic disk storage]**
Friedman, M.B.;
Spectrum, IEEE
Volume 33, Issue 4, April 1996 Page(s):73 - 79
Digital Object Identifier 10.1109/6.486635

AbstractPlus | Full Text: PDF(3092 KB) IEEE JNL
Rights and Permissions

6. **Component-based performance modeling of a storage area network**
Aizikowitz, N.; Glikson, A.; Landau, A.; Mendelson, B.; Sandbank, T.;
Winter Simulation Conference, 2005 Proceedings of the
4-7 Dec. 2005 Page(s):10 pp.
Digital Object Identifier 10.1109/WSC.2005.1574534
AbstractPlus | Full Text: PDF(526 KB) IEEE CNF
Rights and Permissions

7. **Improving Disk Throughput in Data-Intensive Servers**
Carrera, E.V.; Bianchini, R.;
High Performance Computer Architecture, 2004. HPCA-10. Proceedings. 10th Symposium on
14-18 Feb. 2004 Page(s):130 - 130
Digital Object Identifier 10.1109/HPCA.2004.10023
AbstractPlus | Full Text: PDF(248 KB) IEEE CNF
Rights and Permissions

8. **A performance evaluation tool for RAID disk arrays**
Thomasian, A.; Han, C.; Fu, G.; Liu, C.;
Quantitative Evaluation of Systems, 2004. QEST 2004. Proceedings. First Inte Conference on the
27-30 Sept. 2004 Page(s):8 - 17
Digital Object Identifier 10.1109/QEST.2004.1348011
AbstractPlus | Full Text: PDF(607 KB) IEEE CNF
Rights and Permissions

9. **Fault-tolerant distributed mass storage for LHC computing**
Wiebalck, A.; Breuer, P.T.; Lindenstruth, V.; Stinbeck, T.M.;
Cluster Computing and the Grid, 2003. Proceedings. CCGrid 2003. 3rd IEEE/A Symposium on
12-15 May 2003 Page(s):266 - 273
Digital Object Identifier 10.1109/CCGRID.2003.1199377
AbstractPlus | Full Text: PDF(350 KB) IEEE CNF
Rights and Permissions

10. **A unified, low-overhead framework to support continuous profiling and o**
Ming Zhang; Xubin He; Qing Yang;
Performance, Computing, and Communications Conference, 2003. Conferenc the 2003 IEEE International
9-11 April 2003 Page(s):327 - 334
AbstractPlus | Full Text: PDF(751 KB) IEEE CNF
Rights and Permissions

11. **Design and implementation of a block storage multi-protocol converter**
Gerasimov, I.; Zhuravlev, A.; Pershin, M.; Gerasimov, D.V.;
Mass Storage Systems and Technologies, 2003. (MSST 2003). Proceedings. : NASA Goddard Conference on
7-10 April 2003 Page(s):183 - 188
AbstractPlus | Full Text: PDF(285 KB) IEEE CNF
Rights and Permissions

12. **CoStore: a reliable and highly available storage system using clusters**
Yong Chen; Ni, L.M.; Chengzhong Xu; Mingyao Yang; Kusler, J.F.; Pei Zheng;
High Performance Computing Systems and Applications, 2002. Proceedings. International Symposium on
16-19 June 2002 Page(s):3 - 11

Digital Object Identifier 10.1109/HPCSA.2002.1019128

AbstractPlus | Full Text: PDF(345 KB) **IEEE CNF**
Rights and Permissions

☐ **13. Servers in news and sport: evolution and revolution**
Stewart, J.;
Broadcasting Convention, 1997. International
12-16 Sept. 1997 Page(s):252 - 257

AbstractPlus | Full Text: PDF(512 KB) **IEE CNF**

Help Contact Us Privacy & :

PORTAL

USPTO

Subscribe (Full Service) Register (Limited Service, Free) Login

Search: ◉ The ACM Digital Library ○ The Guide

SEARCH

Advanced Search

Search Tips

Enter words, phrases or names below. Surround phrases or full names with double quotation marks.

Desired Results:

must have **all** of the words or phrases

RAID, cache

must have **any** of the words or phrases

configuration, reconfiguration, change RAID level

must have **none** of the words or phrases

Name or Affiliation:

Authored by: ◉ all ○ any ○ none

Edited by: ◉ all ○ any ○ none

Reviewed by: ◉ all ○ any ○ none

Only search in:*

○ Title ◉ Abstract ○ Review ○ All Information

SEARCH

*Searches will be performed on all available information, including full text where available, unless specified above.

ISBN / ISSN: ◉ Exact ○ Expand

DOI: ◉ Exact ○ Expand

SEARCH

Published:

By: ◉ all ○ any ○ none

In: ◉ all ○ any ○ none

Since:

Month Year

Before:

Month Year

As: Any type of publication

Conference Proceeding:

Sponsored By:

Conference Location:

Conference Year:

yyyy

SEARCH

Classification: (CCS) ☐ **Primary Only**

Classified as: ◉ all ○ any ○ none

Subject Descriptor: ◉ all ○ any ○ none

Keyword Assigned: ◉ all ○ any ○ none

Results must have accessible:

☐ Full Text ☐ Abstract ☐ Review

Terms used
RAID cache configuration reconfiguration change RAID level

Found **10** of **178,880**

Sort results by: relevance
Display results: expanded form

Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 10 of 10

Relevance scale

1 The HP AutoRAID hierarchical storage system
John Wilkes, Richard Golding, Carl Staelin, Tim Sullivan
February 1996 **ACM Transactions on Computer Systems (TOCS)**, Volume 14 Issue 1
Publisher: ACM Press
Full text available: pdf(1.82 MB) Additional Information: full citation, abstract, references, citings, index terms

Configuring redundant disk arrays is a black art. To configure an array properly, a system administrator must understand the details of both the array and the workload it will support. Incorrect understanding of either, or changes in the workload over time, can lead to poor performance. We present a solution to this problem: a two-level storage hierarchy implemented inside a single disk-array controller. In the upper level of this hierarchy, two copies of active data are stored to provide f ...

Keywords: RAID, disk array, storage hierarchy

2 EERAID: energy efficient redundant and inexpensive disk array
Dong Li, Jun Wang
September 2004 **Proceedings of the 11th workshop on ACM SIGOPS European workshop: beyond the PC EW11**
Publisher: ACM Press
Full text available: pdf(193.16 KB) Additional Information: full citation, abstract, references

Recent research works have been presented on conserving energy for multi-disk systems either at a single disk drive level or at a storage system level and thereby having certain limitations. This paper studies several new redundancy-based, power-aware, I/O request scheduling and cache management policies at the RAID controller level to build energy-efficient RAID systems, by exploiting the redundant information and destage issues of the array for two popular RAID levels, RAID 1 and RAID 5. For R ...

3 Hierarchical disk cache management in RAID 5 controller
Jung-ho Huh, Tae-mu Chang
December 2003 **Journal of Computing Sciences in Colleges**, Volume 19 Issue 2
Publisher: Consortium for Computing Sciences in Colleges
Full text available: pdf(137.71 KB) Additional Information: full citation, abstract, references, index terms

In RAID system, disk cache is one of the important elements in improving the system

performance. Two-level cache displays superior performance in comparison to single cache and is effective in temporal and spatial locality. The proposed cache system consists in two levels. The first level cache is a set associative cache with small block size whereas the second level cache is a fully associative cache with large block size. In this paper, a RAID 5 disk cache model is presented that is located on ...

4 X-RAY: A Non-Invasive Exclusive Caching Mechanism for RAIDs

Lakshmi N. Bairavasundaram, Muthian Sivathanu, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau

March 2004 **ACM SIGARCH Computer Architecture News , Proceedings of the 31st annual international symposium on Computer architecture ISCA '04**, Volume 32 Issue 2

Publisher: IEEE Computer Society, ACM Press

Full text available: pdf(250.59 KB) Additional Information: full citation, abstract, citings

RAID storage arrays often possess gigabytes of RAM forcaching disk blocks. Currently, most RAID systems use LRUor LRU-like policies to manage these caches. Since these arraycaches do not recognize the presence of file system buffer caches,they redundantly retain many of the same blocks as those cachedby the file system, thereby wasting precious cache space. In thispaper, we introduce X-RAY, an exclusive RAID array cachingmechanism. X-RAY achieves a high degree of (but not perfect) exclusivitythr ...

5 Parity logging overcoming the small write problem in redundant disk arrays

Daniel Stodolsky, Garth Gibson, Mark Holland

May 1993 **ACM SIGARCH Computer Architecture News , Proceedings of the 20th annual international symposium on Computer architecture ISCA '93**, Volume 21 Issue 2

Publisher: ACM Press

Full text available: pdf(1.35 MB) Additional Information: full citation, abstract, references, citings, index terms

Parity encoded redundant disk arrays provide highly reliable, cost effective secondary storage with high performance for read accesses and large write accesses. Their performance on small writes, however, is much worse than mirrored disks—the traditional, highly reliable, but expensive organization for secondary storage. Unfortunately, small writes are a substantial portion of the I/O workload of many important, demanding applications such as on-line transaction processing. This paper ...

6 Parity logging disk arrays

Daniel Stodolsky, Mark Holland, William V. Courtright, Garth A. Gibson

August 1994 **ACM Transactions on Computer Systems (TOCS)**, Volume 12 Issue 3

Publisher: ACM Press

Full text available: pdf(1.98 MB) Additional Information: full citation, abstract, references, citings, index terms

Parity-encoded redundant disk arrays provide highly reliable, cost-effective secondary storage with high performance for reads and large writes. Their performance on small writes, however, is much worse than mirrored disks—the traditional, highly reliable, but expensive organization for secondary storage. Unfortunately, small writes are a substantial portion of the I/O workload of many important, demanding applications such as on-line transaction processing. This paper presents

Keywords: RAID, disk arrays

7 Destage algorithms for disk arrays with non-volatile caches

Anujan Varma, Quinn Jacobson

May 1995 **ACM SIGARCH Computer Architecture News , Proceedings of the 22nd**

annual international symposium on Computer architecture ISCA '95, Volume 23 Issue 2
Publisher: ACM Press

Full text available: pdf(1.63 MB)   Additional Information: full citation, abstract, references, citings, index terms

In a disk array with a nonvolatile write cache, destages from the cache to the disk are performed in the background asynchronously while read requests from the host system are serviced in the foreground. In this paper, we study a number of algorithms for scheduling destages in a RAID-5 system. We introduce a new scheduling algorithm, called *linear threshold scheduling,* that adaptively varies the rate of destages to disks based on the instantaneous occupancy of the write cache. The perform ...

8 The architecture of a fault-tolerant cached RAID controller
Jai Menon, Jim Cortney
May 1993 **ACM SIGARCH Computer Architecture News , Proceedings of the 20th annual international symposium on Computer architecture ISCA '93**, Volume 21 Issue 2
Publisher: ACM Press

Full text available: pdf(1.04 MB)   Additional Information: full citation, abstract, references, citings, index terms

RAID-5 arrays need 4 disk accesses to update a data block—2 to read old data and parity, and 2 to write new data and parity. Schemes previously proposed to improve the update performance of such arrays are the Log-Structured File System [10] and the Floating Parity Approach [6]. Here, we consider a third approach, called Fast Write, which eliminates disk time from the host response time to a write, by using a Non-Volatile Cache in the disk array controller. We examine three alternativ ...

9 Performance of a disk array protype
Ann L. Chervenak, Randy H. Katz
April 1991 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1991 ACM SIGMETRICS conference on Measurement and modeling of computer systems SIGMETRICS '91**, Volume 19 Issue 1
Publisher: ACM Press

Full text available: pdf(1.05 MB)   Additional Information: full citation, abstract, references, citings, index terms

The RAID group at U.C. Berkeley recently built a prototype disk array. This paper examines the performance limits of each component of the array usiug SCSI bus traces, Sprite operating system traces and user programs.The array performs successfully for a workload of small, random I/O operations, achieving 275 I/Os per second on 14 disks before the Sun4/280 host becomes CPU-limited. The prototype is less successful in delivering high throughput for large, sequential operations. Memory system cont ...

10 Proceedings of the IEEE 2001 symposium on parallel and large-data visualization and graphics
David Breen, Alan Heirich, Anton Koning
October 2001 proceeding
Publisher: IEEE Press
Additional Information: full citation, abstract

Welcome to the 2001 IEEE Symposium on Parallel and Large Data Visualization and Graphics (PVG2001).

This year's proceedings are dominated by scientific applications of parallel visualization, a subject that will be explored in the panel discussion. Other noticeable trends this year are the increasing prevalence of PC graphics clusters and the use of out-of-core techniques as

alternatives to traditional parallelism. The following are summary descriptions of the program presentations.

Results 1 - 10 of 10

Subscribe (Full Service) Register (Limited Service, Free) Login

Search: ⦿ The ACM Digital Library ○ The Guide

SEARCH

USPTO

Advanced Search

Search Tips

Enter words, phrases or names below. Surround phrases or full names with double quotation marks.

Desired Results:
must have **all** of the words or phrases
RAID configuration
must have **any** of the words or phrases
reconfiguration, change RAID level
must have **none** of the words or phrases

Name or Affiliation:
Authored by: ⦿ all ○ any ○ none
Edited by: ⦿ all ○ any ○ none
Reviewed by: ⦿ all ○ any ○ none

Only search in:*
○ Title ⦿ Abstract ○ Review ○ All Information

SEARCH

*Searches will be performed on all available information, including full text where available, unless specified above.

ISBN / ISSN: ⦿ Exact ○ Expand

DOI: ⦿ Exact ○ Expand

SEARCH

Published:
By: ⦿ all ○ any ○ none
In: ⦿ all ○ any ○ none
Since:
Month Year
Before:
Month Year
As: Any type of publication

Conference Proceeding:
Sponsored By:
Conference Location:
Conference Year:
yyyy

SEARCH

Classification: (CCS) ☐ **Primary Only**
Classified as: ⦿ all ○ any ○ none
Subject Descriptor: ⦿ all ○ any ○ none
Keyword Assigned: ⦿ all ○ any ○ none

Results must have accessible:
☐ Full Text ☐ Abstract ☐ Review

Terms used
RAID configuration reconfiguration change RAID level

Found 6 of 178,880

Sort results by: relevance
Display results: expanded form
Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 6 of 6

Relevance scale

1 The HP AutoRAID hierarchical storage system
John Wilkes, Richard Golding, Carl Staelin, Tim Sullivan
February 1996 **ACM Transactions on Computer Systems (TOCS)**, Volume 14 Issue 1
Publisher: ACM Press
Full text available: pdf(1.82 MB) Additional Information: full citation, abstract, references, citings, index terms

Configuring redundant disk arrays is a black art. To configure an array properly, a system administrator must understand the details of both the array and the workload it will support. Incorrect understanding of either, or changes in the workload over time, can lead to poor performance. We present a solution to this problem: a two-level storage hierarchy implemented inside a single disk-array controller. In the upper level of this hierarchy, two copies of active data are stored to provide f ...

Keywords: RAID, disk array, storage hierarchy

2 Deconstructing storage arrays
Timothy E. Denehy, John Bent, Florentina I. Popovici, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau
October 2004 **ACM SIGPLAN Notices , ACM SIGARCH Computer Architecture News , ACM SIGOPS Operating Systems Review , Proceedings of the 11th international conference on Architectural support for programming languages and operating systems ASPLOS-XI**, Volume 39 , 32 , 38 Issue 11 , 5 , 5
Publisher: ACM Press
Full text available: pdf(1.74 MB) Additional Information: full citation, abstract, references, index terms

We introduce Shear, a user-level software tool that characterizes RAID storage arrays. Shear employs a set of controlled algorithms combined with statistical techniques to automatically determine the important properties of a RAID system, including the number of disks, chunk size, level of redundancy, and layout scheme. We illustrate the correctness of Shear by running it upon numerous simulated configurations, and then verify its real-world applicability by running Shear on both software-based ...

Keywords: RAID, storage

3

Design, implementation, and performance of storage systems: Multi-level RAID for very large disk arrays
Alexander Thomasian
March 2006 **ACM SIGMETRICS Performance Evaluation Review**, Volume 33 Issue 4
Publisher: ACM Press
Full text available: pdf(1.25 MB) Additional Information: full citation, abstract, references

Very Large Disk Arrays - VLDAs have been developed to cope with the rapid increase in the volume of data generated requiring ultrareliable storage. Bricks or Storage Nodes - SNs holding a dozen or more disks are cost effective VLDA building blocks, since they cost less than traditional disk arrays. We utilize the Multilevel RAID - MRAID paradigm for protecting both SNs and their disks. Each SN is a k-disk-failure-tolerant kDFT array, while replication or *l*-node failure tolerance - *l*

4 De-layered grid storage server
H. Shrikumar
July 2005 **ACM SIGBED Review**, Volume 2 Issue 3
Publisher: ACM Press
Full text available: pdf(750.59 KB) Additional Information: full citation, abstract, references, index terms

Networks have become faster and disks have become fatter at a pace that, despite Moore's law, CPU developments have simply not been been able to keep up with. We present a Grid Storage Server which is capable of scaling up to meet the "terabit-terabyte" demands of very large scale grid computation applications with large data sets.The Grid Storage Server is implemented almost completely in silicon, whether FPGA or ASIC; the fast-path of this server does not use a CPU or von Neumann style (instru ...

5 Design and implementation of a configurable mixed-media file system
Silvano Maffeis
October 1994 **ACM SIGOPS Operating Systems Review**, Volume 28 Issue 4
Publisher: ACM Press
Full text available: pdf(333.42 KB) Additional Information: full citation, abstract, index terms

In this paper we describe the design and implementation of a configurable mixed-media file system. The attribute *configurable* means that a file system serving a specific application area can be realized out of a library of reusable file system classes. The attribute *mixed-media* stands for the file system's ability to integrate different media types (RAM, harddisks, WORM optical disks, CDROMs, tape devices, RAIDs etc.) into a virtual storage, and making applications unaware of this ...

6 Tolerating multiple failures in RAID architectures with optimal storage and uniform declustering
Guillermo A. Alvarez, Walter A. Burkhard, Flaviu Cristian
May 1997 **ACM SIGARCH Computer Architecture News , Proceedings of the 24th annual international symposium on Computer architecture ISCA '97**, Volume 25 Issue 2
Publisher: ACM Press
Full text available: pdf(1.50 MB) Additional Information: full citation, abstract, references, citings, index terms

We present DATUM, a novel method for tolerating multiple disk failures in disk arrays. DATUM is the first known method that can mask any given number of failures, requires an optimal amount of redundant storage space, and spreads reconstruction accesses uniformly over disks in the presence of failures without needing large layout tables in controller memory. Our approach is based on information dispersal, a coding technique that admits an efficient hardware implementation. As t ...

Results 1 - 6 of 6

acm PORTAL
USPTO

Subscribe (Full Service) Register (Limited Service, Free) Login

Search: ◉ The ACM Digital Library ○ The Guide

SEARCH

[redacted]

Advanced Search

Search Tips

Enter words, phrases or names below. Surround phrases or full names with double quotation marks.

Desired Results:

must have **all** of the words or phrases

RAID

must have **any** of the words or phrases

configuration, reconfiguration, change RAID level

must have **none** of the words or phrases

Name or Affiliation:

Authored by: ◉ all ○ any ○ none

Edited by: ◉ all ○ any ○ none

Reviewed by: ◉ all ○ any ○ none

Only search in:*

○ Title ◉ Abstract ○ Review ○ All Information

SEARCH

*Searches will be performed on all available information, including full text where available, unless specified above.

ISBN / ISSN: ◉ Exact ○ Expand

DOI: ◉ Exact ○ Expand

SEARCH

Published:

By: ◉ all ○ any ○ none

In: ◉ all ○ any ○ none

Since:

Month Year

Before:

Month Year

As: Any type of publication

Conference Proceeding:

Sponsored By:

Conference Location:

Conference Year:

yyyy

SEARCH

Classification: (CCS) ☐ **Primary Only**

Classified as: ◉ all ○ any ○ none

Subject Descriptor: ◉ all ○ any ○ none

Keyword Assigned: ◉ all ○ any ○ none

Results must have accessible:

☐ Full Text ☐ Abstract ☐ Review

PORTAL
USPTO

Subscribe (Full Service) Register (Limited Service, Free) Login

Search: The ACM Digital Library The Guide

+abstract:RAID abstract:configuration, abstract:reconfiguratio

SEARCH

THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Terms used
RAID configuration reconfiguration change RAID level

Found **48** of **178,880**

Sort results by: relevance
Display results: expanded form

Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 20 of 48

Result page: **1** 2 3 next

Relevance scale

1 The HP AutoRAID hierarchical storage system
John Wilkes, Richard Golding, Carl Staelin, Tim Sullivan
February 1996 **ACM Transactions on Computer Systems (TOCS)**, Volume 14 Issue 1
Publisher: ACM Press
Full text available: pdf(1.82 MB) Additional Information: full citation, abstract, references, citings, index terms

Configuring redundant disk arrays is a black art. To configure an array properly, a system administrator must understand the details of both the array and the workload it will support. Incorrect understanding of either, or changes in the workload over time, can lead to poor performance. We present a solution to this problem: a two-level storage hierarchy implemented inside a single disk-array controller. In the upper level of this hierarchy, two copies of active data are stored to provide f ...

Keywords: RAID, disk array, storage hierarchy

2 Deconstructing storage arrays
Timothy E. Denehy, John Bent, Florentina I. Popovici, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau
October 2004 **ACM SIGPLAN Notices , ACM SIGARCH Computer Architecture News , ACM SIGOPS Operating Systems Review , Proceedings of the 11th international conference on Architectural support for programming languages and operating systems ASPLOS-XI**, Volume 39 , 32 , 38 Issue 11 , 5 , 5
Publisher: ACM Press
Full text available: pdf(1.74 MB) Additional Information: full citation, abstract, references, index terms

We introduce Shear, a user-level software tool that characterizes RAID storage arrays. Shear employs a set of controlled algorithms combined with statistical techniques to automatically determine the important properties of a RAID system, including the number of disks, chunk size, level of redundancy, and layout scheme. We illustrate the correctness of Shear by running it upon numerous simulated configurations, and then verify its real-world applicability by running Shear on both software-based ...

Keywords: RAID, storage

3

RAID: a robust and adaptable distributed system
Bharat Bhargava, John Dilley, John Riedl
September 1986 **Proceedings of the 2nd workshop on Making distributed systems work**
Publisher: ACM Press
Full text available: pdf(429.83 KB) Additional Information: full citation, abstract, references

There is a need to design distributed systems that are not rigid in their choice of algorithms and that are responsive to faults/failures and performance degradation. To meet this challenge, we formalize and experiment with design principles that allow the implementation of an adaptable distributed system. The strategies for dynamic reconfiguration of the subsystems and determining their impact are being studied via experiments on a prototype system called *RAID* under development at Purdue ...

4 Design, implementation, and performance of storage systems: Multi-level RAID for very large disk arrays
Alexander Thomasian
March 2006 **ACM SIGMETRICS Performance Evaluation Review**, Volume 33 Issue 4
Publisher: ACM Press
Full text available: pdf(1.25 MB) Additional Information: full citation, abstract, references

Very Large Disk Arrays - VLDAs have been developed to cope with the rapid increase in the volume of data generated requiring ultrareliable storage. Bricks or Storage Nodes - SNs holding a dozen or more disks are cost effective VLDA building blocks, since they cost less than traditional disk arrays. We utilize the Multilevel RAID - MRAID paradigm for protecting both SNs and their disks. Each SN is a k-disk-failure-tolerant kDFT array, while replication or *l*-node failure tolerance - *l*

5 Striping in a RAID level 5 disk array
Peter M. Chen, Edward K. Lee
May 1995 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1995 ACM SIGMETRICS joint international conference on Measurement and modeling of computer systems SIGMETRICS '95/PERFORMANCE '95**, Volume 23 Issue 1
Publisher: ACM Press
Full text available: pdf(1.26 MB) Additional Information: full citation, abstract, references, citings, index terms

Redundant disk arrays are an increasingly popular way to improve I/O system performance. Past research has studied how to stripe data in non-redundant (RAID Level 0) disk arrays, but none has yet been done on how to stripe data in redundant disk arrays such as RAID Level 5, or on how the choice of striping unit varies with the number of disks. Using synthetic workloads, we derive simple design rules for striping data in RAID Level 5 disk arrays given varying amounts of workload information. We t ...

6 Reliability and performance of hierarchical RAID with multiple controllers
Sung Hoon Baek, Bong Wan Kim, Eui Joung Joung, Chong Won Park
August 2001 **Proceedings of the twentieth annual ACM symposium on Principles of distributed computing**
Publisher: ACM Press
Full text available: pdf(663.16 KB) Additional Information: full citation, abstract, references, index terms

Redundant arrays of inexpensive disks (RAID) offer fault tolerance against disk failures. However a storage system having more disks suffers from less reliability and performance. A RAID architecture tolerating multiple disk failures shows severe performance degradation in comparison to the RAID Level 5 due to the complexity of implementation. We present a new RAID architecture that tolerates at least three disk failures and offers similar throughout to the RAID Level 5. We call it the hierar ...

Keywords: Markov process, hierarchical RAID, high reliability, three-failure-tolerant array

7 EERAID: energy efficient redundant and inexpensive disk array

Dong Li, Jun Wang

September 2004 **Proceedings of the 11th workshop on ACM SIGOPS European workshop: beyond the PC EW11**

Publisher: ACM Press

Full text available: pdf(193.16 KB) Additional Information: full citation, abstract, references

Recent research works have been presented on conserving energy for multi-disk systems either at a single disk drive level or at a storage system level and thereby having certain limitations. This paper studies several new redundancy-based, power-aware, I/O request scheduling and cache management policies at the RAID controller level to build energy-efficient RAID systems, by exploiting the redundant information and destage issues of the array for two popular RAID levels, RAID 1 and RAID 5. For R ...

8 Tutorial on storage technology: RAID and beyond

Garth A. Gibson

May 1995 **ACM SIGMOD Record , Proceedings of the 1995 ACM SIGMOD international conference on Management of data SIGMOD '95**, Volume 24 Issue 2

Publisher: ACM Press

Full text available: pdf(104.48 KB) Additional Information: full citation, abstract, references, index terms

In stark contrast to the 25% per year increase in areal density delivered by the magnetic disk industry during the 1970s and 1980s, yearly increases today are 60%, on par with DRAM density increases. Moreover, the storage industry is also delivering substantially higher data rates, smart disk-embedded readahead and writebehind, and a new generation of high-speed serial interconnects. This industry has also embraced Redundant Arrays of Inexpensive (or Independent) Disks (RAID) technology - 1997's ...

9 Hierarchical disk cache management in RAID 5 controller

Jung-ho Huh, Tae-mu Chang

December 2003 **Journal of Computing Sciences in Colleges**, Volume 19 Issue 2

Publisher: Consortium for Computing Sciences in Colleges

Full text available: pdf(137.71 KB) Additional Information: full citation, abstract, references, index terms

In RAID system, disk cache is one of the important elements in improving the system performance. Two-level cache displays superior performance in comparison to single cache and is effective in temporal and spatial locality. The proposed cache system consists in two levels. The first level cache is a set associative cache with small block size whereas the second level cache is a fully associative cache with large block size. In this paper, a RAID 5 disk cache model is presented that is located on ...

10 TRAP-Array: A Disk Array Architecture Providing Timely Recovery to Any Point-in-time

Qing Yang, Weijun Xiao, Jin Ren

June 2006 **Proceedings of the 33rd International Symposium on Computer Architecture (ISCA'06) - Volume 00 ISCA '06**

Publisher: IEEE Computer Society

Full text available: Publisher Site Additional Information: full citation, abstract

RAID architectures have been used for more than two decades to recover data upon disk failures. Disk failure is just one of the many causes of damaged data. Data can be

damaged by virus attacks, user errors, defective software/firmware, hardware faults, and site failures. The risk of these types of data damage is far greater than disk failure with today's mature disk technology and networked information services. It has therefore become increasingly important for today's disk array to be able to ...

11 A case for redundant arrays of inexpensive disks (RAID)

David A. Patterson, Garth Gibson, Randy H. Katz

June 1988 **ACM SIGMOD Record , Proceedings of the 1988 ACM SIGMOD international conference on Management of data SIGMOD '88**, Volume 17 Issue 3

Publisher: ACM Press

Full text available: pdf(1.20 MB) Additional Information: full citation, abstract, references, citings, index terms

Increasing performance of CPUs and memories will be squandered if not matched by a similar performance increase in I/O. While the capacity of Single Large Expensive Disks (SLED) has grown rapidly, the performance improvement of SLED has been modest. Redundant Arrays of Inexpensive Disks (RAID), based on the magnetic disk technology developed for personal computers, offers an attractive alternative to SLED, promising improvements of an order of magnitude in performance, reliability, ...

12 TECH spyware raid: the PEA matrix

Joni Mason, Marilee Tuomanen

November 2005 **Proceedings of the 33rd annual ACM SIGUCCS conference on User services SIGUCCS '05**

Publisher: ACM Press

Full text available: pdf(143.08 KB) Additional Information: full citation, abstract, index terms

On May 1, 2005, the Phillips Exeter Academy Information Technology Services Department (ITS), in conjunction with Student TECHs (Technical Ethernet Crisis Helpers), held its first TECH spyware raid.The theme of the spyware raid was centered on a short film the TECHs wrote and produced, a parody of the popular *Matrix* trilogy called *The PEA Matrix*. In *The PEA Matrix*, student TECHs fight the evil agents of spyware.The screening of the film of was preceded by a media campaign bu ...

Keywords: adware, education, residential, spyware

13 X-RAY: A Non-Invasive Exclusive Caching Mechanism for RAIDs

Lakshmi N. Bairavasundaram, Muthian Sivathanu, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau

March 2004 **ACM SIGARCH Computer Architecture News , Proceedings of the 31st annual international symposium on Computer architecture ISCA '04**, Volume 32 Issue 2

Publisher: IEEE Computer Society, ACM Press

Full text available: pdf(250.59 KB) Additional Information: full citation, abstract, citings

RAID storage arrays often possess gigabytes of RAM forcaching disk blocks. Currently, most RAID systems use LRUor LRU-like policies to manage these caches. Since these arraycaches do not recognize the presence of file system buffer caches,they redundantly retain many of the same blocks as those cachedby the file system, thereby wasting precious cache space. In thispaper, we introduce X-RAY, an exclusive RAID array cachingmechanism. X-RAY achieves a high degree of (but not perfect) exclusivitythr ...

14 EVENODD: an optimal scheme for tolerating double disk failures in RAID architectures

M. Blaum, J. Brady, J. Bruck, J. Menon

April 1994 **ACM SIGARCH Computer Architecture News , Proceedings of the 21ST**

annual international symposium on Computer architecture ISCA '94, Volume 22 Issue 2
Publisher: IEEE Computer Society Press, ACM Press
Full text available: pdf(893.35 KB) Additional Information: full citation, abstract, references, citings, index terms

We present a novel method, that we call EVENODD, for tolerating up to two disk failures in RAID architectures. EVENODD is the first known scheme for tolerating double disk failures that is optimal with regard to both storage and performance. EVENODD employs the addition of only two redundant disks and consists of simple exclusive-OR computations. A major advantage of EVENODD is that it only requires parity hardware, which is typically present in standard RAID-5 controllers. Hence, EVENODD can be ...

15 Massive arrays of idle disks for storage archives
Dennis Colarelli, Dirk Grunwald
November 2002 **Proceedings of the 2002 ACM/IEEE conference on Supercomputing**
Publisher: IEEE Computer Society Press
Full text available: pdf(751.87 KB) Additional Information: full citation, abstract, references, citings, index terms

The declining costs of commodity disk drives is rapidly changing the economics of deploying large amounts of online or near-line storage. Conventional mass storage systems use either high performance RAID clusters, automated tape libraries or a combination of tape and disk. In this paper, we analyze an alternative design using *massive arrays of idle disks*, or MAID. We argue that this storage organization provides storage densities matching or exceeding those of tape libraries with perform ...

16 Reliability and security of RAID storage systems and D2D archives using SATA disk drives
Gordon F. Hughes, Joseph F. Murray
February 2005 **ACM Transactions on Storage (TOS)**, Volume 1 Issue 1
Publisher: ACM Press
Full text available: pdf(94.82 KB) Additional Information: full citation, abstract, references, index terms

Information storage reliability and security is addressed by using personal computer disk drives in enterprise-class nearline and archival storage systems. The low cost of these serial ATA (SATA) PC drives is a tradeoff against drive reliability design and demonstration test levels, which are higher in the more expensive SCSI and Fibre Channel drives. This article discusses the tradeoff between SATA which has the advantage that fewer higher capacity drives are needed for a given system storage c ...

Keywords: Disk drive, SATA, SMART, archival storage, failure prediction, secure erase, storage resource management, storage systems architecture

17 Parity logging overcoming the small write problem in redundant disk arrays
Daniel Stodolsky, Garth Gibson, Mark Holland
May 1993 **ACM SIGARCH Computer Architecture News , Proceedings of the 20th annual international symposium on Computer architecture ISCA '93**, Volume 21 Issue 2
Publisher: ACM Press
Full text available: pdf(1.35 MB) Additional Information: full citation, abstract, references, citings, index terms

Parity encoded redundant disk arrays provide highly reliable, cost effective secondary storage with high performance for read accesses and large write accesses. Their performance on small writes, however, is much worse than mirrored disks—the traditional, highly reliable, but expensive organization for secondary storage.

Unfortunately, small writes are a substantial portion of the I/O workload of many important, demanding applications such as on-line transaction processing. This paper ...

18 RAID: high-performance, reliable secondary storage
Peter M. Chen, Edward K. Lee, Garth A. Gibson, Randy H. Katz, David A. Patterson
June 1994 **ACM Computing Surveys (CSUR)**, Volume 26 Issue 2
Publisher: ACM Press
Full text available: pdf(3.60 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Disk arrays were proposed in the 1980s as a way to use parallelism between multiple disks to improve aggregate I/O performance. Today they appear in the product lines of most major computer manufacturers. This article gives a comprehensive overview of disk arrays and provides a framework in which to organize current and future work. First, the article introduces disk technology and reviews the driving forces that have popularized disk arrays: performance and reliability. It discusses the tw ...

Keywords: RAID, disk array, parallel I/O, redundancy, storage, striping

19 Parity logging disk arrays
Daniel Stodolsky, Mark Holland, William V. Courtright, Garth A. Gibson
August 1994 **ACM Transactions on Computer Systems (TOCS)**, Volume 12 Issue 3
Publisher: ACM Press
Full text available: pdf(1.98 MB) Additional Information: full citation, abstract, references, citings, index terms

Parity-encoded redundant disk arrays provide highly reliable, cost-effective secondary storage with high performance for reads and large writes. Their performance on small writes, however, is much worse than mirrored disks—the traditional, highly reliable, but expensive organization for secondary storage. Unfortunately, small writes are a substantial portion of the I/O workload of many important, demanding applications such as on-line transaction processing. This paper presents
Keywords: RAID, disk arrays

20 De-layered grid storage server
H. Shrikumar
July 2005 **ACM SIGBED Review**, Volume 2 Issue 3
Publisher: ACM Press
Full text available: pdf(750.59 KB) Additional Information: full citation, abstract, references, index terms

Networks have become faster and disks have become fatter at a pace that, despite Moore's law, CPU developments have simply not been been able to keep up with. We present a Grid Storage Server which is capable of scaling up to meet the "terabit-terabyte" demands of very large scale grid computation applications with large data sets.The Grid Storage Server is implemented almost completely in silicon, whether FPGA or ASIC; the fast-path of this server does not use a CPU or von Neumann style (instru ...

Results 1 - 20 of 48 Result page: **1** 2 3 next